BEST AVAILABLE COPY

1874.US

3Com Invention Disclosure Form

3Com Confidential

IDF Docket No.

This is an Information Disclosure Form for 3Com Corporation and all of its related companies, hereafter referred to as "COMPANY."

Instructions:

- Complete all sections. Attach supplemental pages and other materials.
- Provide enough information to understand your invention.
- Any questions regarding this form or how to fill it out, refer to Kim Clinger (408 326-1754) or Bill Becker (408 764-5485) by phone or e-mail.
- **Return Completed Form To:** Legal Department - M.S. 1308
  5400 Bayfront Plaza
  Santa Clara, California 95052-8145

3Com Corporation
Legal Dept.
Date Stamp:
RECEIVED

1. Title of the Invention | Write a short descriptive title, avoiding coined terms and project names.

IPsec With Network Address Translation

2. Inventors Include names of all persons who made a contribution to the conception and/or reduction to practice of the invention. If there are more than five inventors, use a Supplemental Sheet.

Full Name:	David Grabelsky	M.S. & Location	N.108 Unet
Home Address:	3800 Lee Street, Skokie, IL 60076		
Home Phone:	(847) 679-8482	Work Phone:	(847) 222-2483
Employee No.	19630	Division:	R&D
Supervisor:	Ikhlaq Sidhu	Department:	Advanced Technologies
Citizenship:	USA	:	
Full Name:	Michael Borella	M.S. & Location	N.108 Unet
Home Address:	1208 Haverhill Circle, Naperville, IL 60563		
Home Phone:	(847) 961-3750	Work Phone:	(847) 342-6093
Employee No.	21032	Division:	R&D
Supervisor:	Ikhlaq Sidhu	Department:	Advanced Technologies
Citizenship:	USA	:	
Full Name:	Ikhlaq Sidhu	M.S. & Location	N.105 Unet
Home Address:	403 River Grove Lane, Vernon Hills, IL 60061		
Home Phone:	(847) 634-0610	Work Phone:	(847) 222-2487
Employee No.	15146	Division:	R&D
Supervisor:	Dan Schoo	Department:	Advanced Technologies
Citizenship:	USA	:	
Full Name:	Danny Nessett	M.S. & Location	3219 Santa Clara
Home Address:	34810 Wabash River Place, Fremont, CA 94555		
Home Phone:	(510) 745-7543	Work Phone:	(408) 326-1169
Employee No.	05123	Division:	Technology Development Center
Supervisor:	Elaine Lusher	Department:	CTO
Citizenship:	USA	:	
Full Name:		M.S. & Location	
Home Address:			
Home Phone:		Work Phone:	
Employee No.		Division:	
Supervisor:		Department:	
Citizenship:		:	

Check here and attach supplemental sheet if more than 5 inventors ☐

inventors' initials DG WB FS

3Com Invention Disclosure Form

3Com Confidential
IDF Docket No.

3. Conception of the Invention

Date of conception: _____ Location of conception **Unet**

Date of first written description: _____

Location of such first description **Unet** Page(s): **7**

Please attach copies of all pertinent lab notebooks or equivalent. They should be signed, dated and witnessed by two other people who have read and understood the description of the invention.

4. Reduction to Practice

Reduction to practice is demonstrating that an invention merely works with evidence such as working models, prototypes or simulations. Reduction to practice is not necessary to file a patent application.

Date of any such demonstration: none planned ☒

date/planned date of demonstration _____

Location of demonstration: _____

5. Invention Applicability/Project/Release/Sale Information

A sale, offer for sale, public showing or release of the invention may affect COMPANY's right to patent the invention. Submit this form even if a public sale, showing or release has occurred.

To which division or operation does this invention best apply?	**R&D**
Field of technology (e.g., manufacturing, switches, hubs, routers, network management software, adapters, etc.)	**Routers, swithes, security**
Project name and description:	
Product name and model number:	
Does this invention relate to an actual or proposed standard or defacto standard?	☐ NO ☒ YES -(name of std.) Distributed Network Address Translation
Please list docket numbers of all other invention disclosures that are related to this one:	**1572, 1868**
Estimated/actual date of first public release or showing of invention or product incorporating or using the invention:	
Estimated/actual manufacturing release date of invention or product incorporating or using the invention:	

inventors' initials DG MB FS [illegible] ____

3Com Invention Disclosure Form

3Com Confidential
IDF Docket No.

Estimated/actual date of offer for sale of product incorporating or using the invention:	

6. Publication of the Invention

Publication of a description of the invention may affect COMPANY's right to patent the invention. Submit this form even if publication has occurred.

Has a description been published or is it scheduled to be published?

If "Yes," when and to whom? _____

Has a description been disclosed or is it scheduled to be disclosed?

If "Yes," when and to whom? _____

Was a Non-Disclosure Agreement used? No ☐ Yes ☐

If "Yes," please attach copy. _____

7. Government Agency Contract

Was this invention made under a government agency contrac

If "Yes," government agency contract number: _____

8. Joint Development or Development Contract

Was this invention jointly developed with inventors from another compar

If "Yes," please identify the company and/or non COMPANY inventors: _____

Was the invention tested, constructed or conceived pursuant to the performance of a development contract with another company?

If "Yes," please identify the contract and its location: _____

9. Related Art

Is this invention an improvement of an existing COMPANY product?

If "Yes," identify the existing product: _____

Closest known related art ("prior art"): _____

What was the problem to be solved? _____

How had others attempted to solved it before you? _____

What were the problems or disadvantages with prior solutions? _____

10. Key Words

Indicate any key words (preferably at least 3) we could use to search **IPsec, Network Address Translation,**

inventors' initials DH MG IS Aboy

for related art or identify this invention for use in our own database.

Distributed Network Address Translation

11. Reference Materials

List any printed publications, patents, patent applications or any other materials you are aware of which provides background material and/or prior art for your invention

3Com patents on Distributed Network Address Translation (DNAT), and Integrating Network Address Translation with Mobile IP by same authors; Internet draft on DNAT; additional references in supplemental sheets

12. Brief Description of the Invention

Describe the structure, function and/or method of the invention in just enough detail to enable someone technical to understand your invention. Stress the fundamental principle of the new idea from an engineering standpoint. Attach all the relevant descriptive materials. You should reference/include any drawings/sketches that will help explain the invention.

3Com Invention Disclosure Form

3Com Confidential

IDF Docket No.

This invention disclosure describes a method for supporting IPsec between two hosts where one or both of them reside on a stub network behind a router that does network address translation. Specifically, this method applies to a stub network that uses Distributed Network Address Translation (DNAT). The hosts on the stub network are assumed to have only locally routable IP addresses, and must rely on the DNAT router for any global communications on the Internet. In the absence of IPsec, global routing for hosts on the stub networks is accomplished using address translation between port numbers and local IP addresses of the hosts. The new method described here enables IPsec to operated with DNAT by using two main elements: 1) address translation is accomplished by a mapping between the (local) IP address of a given local host and the security parameter index (SPI) associated with each inbound security association (SA) that terminates at the host; and 2) the DNAT router is configured to act as a local certificate authority that may vouch for the identities of hosts on its stub network, allowing hosts to bind a public key to a name space which combines the DNAT router's global IP address and a unique set of port numbers. The DNAT router issues the certificates, and may itself be authenticated with a higher certificate authority. Using a local certificate, any local host on the stub network may initiate and be the termination point of an SA to any other host (or security gateway) on the Internet. The intent of this method is to allow end-to-end SA establishment and operation across DNAT routers. This method will also work when DNAT is used with Mobile IP.

13. Drawings of the Invention

Please submit clear drawings and/or sketches which illustrate the invention either by electronically inserting them into Section 12 or using the supplemental sheets if you can't easily electronically insert them.

14. Licensing/Competitive Analysis

Is this a licensable technology? No ☐ Yes ☒ Don't know ☐

If "Yes," name the fields in which this might be licensed: **Routers, swithes, security**

If "Yes," name the companies which may possibly be interested: _____

15. Attorney

If there is a particular patent attorney with whom you would like to work on this disclosure, suggest his/her name. **Steve Lesavich**

16. Signatures of Inventors

This Invention Disclosure Form is submitted pursuant to your employment agreement with COMPANY. Use a Supplemental Sheet if there are more than 5 inventors. Please sign and date below and be certain that each page of this disclosure has been initialed by each inventor.

inventors' Initials PH MB IS [illegible] ____

3Com Invention Disclosure Form

3Com Confidential
IDF Docket No.

Signature: _____ Date: _____

Signature: _____ Date: _____

Signature: _____ Date: _____

Signature: _____ Date: _____

Signature: _____ Date: _____

7. Witnesses Read and Understood

This Invention Disclosure Form consisting of pages has been read and understood by:

Name: Jerry Mahler Name: Jacek Grabiec

Signature: _____ Signature: _____

Date: _____ Date: _____

Witnesses, please initial all supplemental pages.

inventors' initials DY MB IS Dan ____

1894.
CS

3Com Invention Disclosure Form

3Com Confidential
IDF Docket No.

This is a SUPPLEMENTAL INVENTOR SHEET which is to be used if there are more than 5 inventors for the invention set out in the Invention Docket specified above.

Additional Inventors Include names of all persons who made a contribution to the conception and/or reduction to practice of the invention. If there are more than five inventors, use a Supplemental Sheet.

Full Name:
Home Address: M.S. & Location
Home Phone: Work Phone:
Employee No. Division:
Supervisor: Department:
Citizenship: :

Full Name:
Home Address: M.S. & Location
Home Phone: Work Phone:
Employee No. Division:
Supervisor: Department:
Citizenship: :

Full Name:
Home Address: M.S. & Location
Home Phone: Work Phone:
Employee No. Division:
Supervisor: Department:
Citizenship: :

Full Name:
Home Address: M.S. & Location
Home Phone: Work Phone:
Employee No. Division:
Supervisor: Department:
Citizenship: :

Full Name:
Home Address: M.S. & Location
Home Phone: Work Phone:
Employee No. Division:
Supervisor: Department:
Citizenship: :

inventors' initials DY MB IS DTQ ____

3Com Invention Disclosure Form

3Com Confidential
IDF Docket No.____

This is a SUPPLEMENTAL INFORMATION SHEET to be used to provide additional information regarding the invention disclosure referenced above.

Item ___

Disclosure:
IPsec With Network Address Translation

David Grabelsky
Mike Borella
Ikhlaq Sidhu
Advanced Technologies Department
Carrier Systems Business Unit, 3Com
Dan Nessett
Technology Development Center, 3Com

Original Conception Date: 30 March, 1998

Introduction

IPsec is a protocol for implementing security for Internet communications at the IP layer. Communications between two networking devices or computers over an IP traffic flow made secure by IPsec are afforded the protections of IPsec without having to implement application layer security. The fundamental principle of IPsec is to provide connections that are end-to-end secure on an IP packet-by-packet basis. Only the IPsec entities at the connection endpoints have access to, and participate in, the critical and sensitive operations that make their common connection secure, and each endpoint entity has confidence in the security of the connection and the authenticity of their counterpart at the other end.

Network Address Translation, or NAT (cf., [1]), is a method for bridging networks that use different address spaces. The functional block or device that performs the translation is usually referred to simply as the NAT. In the context of IP networks, different address spaces could be local versus global addresses, or IPv4 versus IPv6 addresses. In the former case, the local address space may be used to alleviate the scarcity of global IP (v4) addresses by placing a stub network (e.g., a LAN) behind a single globally routable gateway/router device (the NAT). Hosts on the stub network possess IP addresses which are only locally routable (i.e., within the stub network). Access to the global IP address space is provided through the NAT. In the latter case, the NAT is required to support the coexistence of IPv4 and IPv6 networks during the potentially long time required for the transition from IPv4 to IPv6. In either case, standard NAT as defined in, e.g., [1] violates certain specific principles of IPsec which are at the foundation of establishment and maintenance of secure end-to-end connections over the IP network.

This disclosure describes a proposed method for enabling IPsec between hosts when one or both of them sit on a stub network behind a NAT which uses a mapping between TCP or UDP port number and local IP address. In this case NAT is used to alleviate the address scarcity problem. Furthermore, we consider a baseline configuration which uses a distributed form of NAT, called DNAT, previously described and disclosed by several of the authors [2]. The basic idea is to configure the DNAT router to authenticate local hosts on its stub network, and to use a mapping between an IPsec-specific identifier

inventors' initials DG MB ID [illegible] ___
JM JG

and local IP address. As described below, these methods resolve the chief objections to implementing IPsec within a NAT environment.

Below, we first provide a very brief summary of IPsec, distilling only certain fundamental principles upon which it is based, and which are violated by standard NAT. Next, we give a brief description of NAT, and how it conflicts with IPsec. This is followed by a summary of DNAT. Finally, we present our proposed method for making IPsec operational in concert with DNAT.

IPsec: The Essentials

IPsec provides security on an IP packet-by-packet basis between two endpoints that terminate a logical IP network connection. The endpoints of the connection are either host computers or security gateways. Here, a host computer is taken to be one with a single network interface; i.e., no routing capabilities. A security gateway is a computer with multiple network interfaces, and routing capabilities between them. In either case, the IPsec protocols are implemented and operational on each endpoint of the connection.

Detailed descriptions of IPsec and related protocols can be found in the references [3-8].

For the purposes of this disclosure, we focus only on two fundamental requirements which are addressed by IPsec:

1. Provide message authentication, integrity and confidentiality services for IP packets moving between a source and destination system.

2. Starting from a state in which no connection exists between two endpoints, establish a security association (based upon IP) between them with the properties that: i) each endpoint trusts the security of the connection; and ii) each identity of each endpoint is authenticated to the other.

These two requirements do not nearly sum up IPsec, but by describing in somewhat abstract terms the approaches taken to address them, we can recognize the difficulties in implementing IPsec with NAT.

To address (1), IPsec defines two security services, each having an associated header that is added to the IP packet that it protects. These are the Authentication Header (AH) and Encapsulating Security Payload (ESP) header. AH provides authentication and integrity protection for each IP packet; ESP provides encryption protection, as well as optional authentication and integrity protection. These IPsec protocol headers are identified in the protocol field of the packet's IP packet header. The IPsec protocol header specifies the protocol type (AH or ESP), and contains a numerical value called the Security Parameter Index (SPI), which is a unique identifier associated by the receiving system with the security association (the header also contains other fields not described here). The identifying information is used by the receiving system to help it correctly process the packet, e.g., to decrypt it or verify its integrity and authenticity.

Each IPsec security service can be applied in one of two modes: transport mode, or tunnel mode. In transport mode, each packet is routed directly to its final destination according to its destination address; the final destination is both where the IPsec receive processing is done, as well as where the packet is consumed. The destination IP address is in the clear as the packet traverses the network. In tunnel mode, the outermost IP header encapsulates the protected IP packet; the final destination of the packet is not necessarily the same as the endpoint of the tunnel, and the address of the final destination in the IP header of the encapsulated (protected) packet may or may not be in the clear.

The IPsec protocols establish and use a Security Association, or SA, to identify a secure "channel" between two endpoints. An SA is a unidirectional connection between these systems that represents a single IPsec protocol-mode combination. The two termination points (end systems in the case of transport services or intermediate devices in the case of tunnel mode services) of a single SA define the logical connection that is protected by IPsec services; one of the endpoints sends IP

packets, the other receives them. Since an SA is unidirectional, a minimum of two SAs is required for secure, bi-directional communications. It is also possible to configure multiple layers of IPsec protocols between two endpoints by combining multiple SAs.

An IPsec implementation includes protocols for managing and processing SAs, and thereby the machinery for implementing the IPsec protocols. For IP packet transmission, the protocols specify the type and order of the processing required based upon information contained in the packet. A single packet may be associated with multiple SAs, whereby a sender applies multiple-SA protection iteratively, one SA after another. For a given IP packet, the specific combination of SAs applied is determined according to parameters carried in the packet's IP header, IPsec header, and/or TCP or UDP header. Examples of these parameters, called selectors in the context of IPsec, include destination IP address, and TCP or UDP port number. Output IPsec processing uses the selectors as a sort of IP packet filter to control the processing options on a packet-by-packet basis. Additional SAs may also be applied by other computers along the path to the final destination. The result is IPsec protection by one or more levels of SAs. For IP packet reception, the protocols specify how to "undo" (i.e., decrypt and/or authenticate) the protection for each level of SA that terminates at the receiving end. Each SA that terminates at a receiving IPsec entity is uniquely identified by the combination of protocol type (AH or ESP), destination IP address, and SPI. These three pieces of information are carried in every incoming packet associated with an SA, and enable the receiving system to determine each SA associated with a given packet by examination. Knowing the associated SA(s) allows the receiving computer to process the SA protection applied by the sending computer(s). If the system is a security gateway, the result of input processing may be an IP packet that is forwarded to a subsequent destination where further IPsec receive processing is done. Further details of IPsec are found in the references mentioned above.

To address (2), a set of protocols has been developed to allow two systems to establish one or more SAs between them. The process of establishing an IPsec SA involves both negotiation and authentication. The negotiation results in an agreement between the two systems as to which security protocol and mode to use, as well as the specific algorithms and associated parameter values, and SPI assignment. The authentication ensures that each end system can trust the identity of its counterpart during negotiation, and hence the once the SA is established.

A number of standards have been proposed for the protocols that accomplish (2), among them the Internet Security Association and Key Exchange Protocol (ISAKMP [6]), and the Oakley protocol [7]. Recently, the Internet Engineering Task Force released the Internet Key Exchange (IKE) protocol [8], which incorporates Oakley into ISAKMP. The operation of ISAKMP and IKE are summarized as follows. Negotiation is carried out as a sequence of signaling exchanges between the two endpoints, in which an initiator proposes the protocol and algorithms, and the responder accepts or counter-proposes. Once the signaling is complete, both endpoints have agreed to the details, exchanged relevant security parameter information, and are ready to send/receive on a single SA. Authentication is based upon a trusted third-party called the Certificate Authority, or CA. Each system (host or security gateway) that participates in IPsec generates a public/private key pair, and has its public key "notarized" by the CA. The CA binds the system's IP address to the its public key, generates a certificate, and returns it to the owner of the key; thus, IP addresses are the name space for binding public keys to their owners. At some point during SA negotiation, each party supplies the other with its certificate, along with a signature that is encrypted with its private key, and that can only be verified with its public key. The recipient (one at each endpoint) uses sender's public key from its certificate to validate the signature and the sender's right to use its IP address. Since only the initiator (sender) has access to the private key, the recipient, once it has verified the signature, is certain of the initiator's identity. Here, "identity" refers only to the IP address of the initiator, as IP addresses form the name space used to bind public keys to their owners. Note that certificates are issued with a lifetime, after which they expire and become invalid. The result of these negotiation and authentication procedures is a secure connection that satisfies the properties of requirement (2) above.

NAT and Why It Doesn't Work With IPsec

Network Address Translation (NAT) allows a single, globally routable IP address to be shared among multiple networked local host computers, none of which has a globally routable IP address. Each local host has a local-only IP address for

communication with its peers on the LAN. A single gateway computer with one local (internal) interface to the LAN and another global (external) interface provides the shared, globally routable IP address. Such a configuration is called a stub network. The gateway routes packets between the global IP network and the local IP network of the LAN, using NAT to map packets between the local and global address spaces that it bridges. The gateway is referred to as the NAT router.

NAT uses a mapping between local IP addresses and TCP or UDP port number to correctly route packets between local hosts and computers on the global IP network. The basic approach is as follows.

When a local host has an IP packet to send, e.g., to a server on the global IP network, it sets its own local IP address as the source and the global IP address of the server as the destination. The NAT router recognizes the packet as bound for the global IP network by its destination address, and replaces the source IP address with that of its external interface. At the same time it records the source TCP or UDP port number of the packet, and associates the port number with the local IP address of the source host; it may also need to change the source port number in the outgoing packet. This state information is maintained at the NAT router for all communications between local hosts and external computers. In some cases, for example FTP connection setup, the IP source address of the local host is also contained in the packet payload. In these cases, the NAT router must additionally modify the packet payload to contain its external IP address, as well as possibly modify the packet length field in the IP header. The IP packet can then be forwarded on the external, global IP network.

When a packet is received from the external network, the NAT router uses the destination TCP or UDP port number to determine which local host should get the packet. Prior to forwarding it to the local host, the destination IP address must be changed to that of the local host. Payload modifications similar to those done for transmission to the external network may also be required.

In very simplified terms, there are two fundamental reasons why NAT doesn't work with IPsec.

First, the NAT router needs to modify the IP packet. However, once an IP packet is protected by IPsec, it cannot be modified anywhere along its path to the IPsec destination. Clearly, the NAT router violates this condition. Even if the NAT router did not need to modify the packets it forwards, it must be able to read the TCP or UDP port number. If ESP is used by the local host, these numbers will be encrypted, so the NAT router won't be able to do its required mapping.

Second, local hosts on a LAN which uses NAT possess only local, non-unique IP addresses. These do not comprise a name space that is suitable for binding a public key to a unique identity. Without this unique binding, it is not possible to provide the authentication necessary for establishment of SAs. Without authentication, neither endpoint can be certain of the identity of their counterpart, and thus cannot establish a secure and trusted connection.

Summary of DNAT

As with standard NAT, DNAT allows a single routable IP address to be multiplexed among several hosts on a local stub network, none of which has a globally routable IP address. However, unlike standard NAT, DNAT allows the router which performs the address mapping to do so without modifying the contents of the routed packets (i.e., TCP/UDP header, or payload). This not only results in a substantial reduction in processing and state-maintenance requirements, but it also eliminates the first of the chief objections to implementing a NAT in the path of an SA: the requirement that the NAT router modify all IP packets that it forwards. The remainder of this section is an abbreviated description of how DNAT works; a more complete description can be found in [2].

Each local host on the stub network has an IP address which supports only local routing; i.e., between all other hosts on the local stub network. The stub network sits behind a single router which has one interface to the local network, and another to the external, global internet. The router's interface to the external networks is via a globally routable IP address, and provides the only access to the external network (definition of a stub network). For local-only communications among hosts and the router, the local IP addresses are used. For communications between a local host and the external, global internet,

DNAT distributes the tasks involved in network address translation between the router and the local host that wants the external communications. A "light-weight tunnel" is established between the local host and the router, and an address mapping between local host IP address and port number is used by the router to route packets between local hosts and the external internet.

The router (hereafter, DNAT router) allocates to each of the local hosts blocks of non-overlapping port numbers. Hosts may also request additional port numbers. The DNAT router maintains a table that maps port numbers to local host IP addresses, and must ensure that no port number is ever allocated to more than one single host at any given time. When transmitting a packet bound for the external internet, the local host sets the source port in the TCP/UDP header (and in the payload, if applicable) to one of the port numbers allocated to it by the DNAT router, and sets the source address in the IP packet header to the *router's* global IP address; the destination IP address is set to that of the external destination. Then the host prepends an additional IP header with local routing information only. Its own local IP address is the source, and the DNAT router's local IP address is the destination. The packet is then transmitted and the additional, local IP header causes the encapsulated global IP packet to be tunneled to the DNAT router (local interface). The DNAT router removes the outer (local) IP header, then forwards the remaining IP packet to the external internet.

When a packet arrives at the DNAT router from the external internet, the router reads the destination port number in the TCP/UDP header, then constructs a local IP header based upon the mapping between port numbers and local host IP addresses. This local IP header provides a tunnel back to the local host. When the local host receives the packet, it removes the local header, then completes packet processing as if it owns the global IP address. It is, of course, assumed here that the source of the packet (e.g., the remote server or host) has set the destination TCP/UDP port number to the source port number in the packet received from the local host (via the DNAT router). The validity of this assumption is briefly addressed below (see item 5 in the list below).

The fundamental concept of DNAT is that, for global communications, each local host acts as if it owns the global IP address of the DNAT router, so that IP packets and payloads for transmission can be constructed by the hosts "ready-to-go" on the external internet, while received packets require no special processing by hosts. This allows the DNAT router to forward packets (in either direction) without the need to modify any of their contents. (For DNAT without IPsec, the DNAT router does need to access information in TCP/UDP headers. This may be opaque in the case of IPsec, but the methods described later in this disclosure show how to solve this problem.) The key elements of DNAT are non-overlapping blocks of port numbers (as allocated by the DNAT router), port number – local IP address mapping (maintained by the DNAT router), and local host participation in the DNAT protocol (requiring modification to the protocol stack).

This brief summary has glossed over a number of details which are addressed in [2]. The following list is meant to simply call attention to these, without further discussion. None of them pose any significant difficulties to DNAT; rather, they are either implied requirements, or implementation options.

1. For "pure" DNAT, the exact placement of the required software in the host's and router's protocol stacks is not absolutely fixed. However in order to support IPsec, the host's and router's DNAT processing must reside above the IP layer.
2. The host implementation will require an appropriate means for determining which outbound packets are destined for the local network, and which ones are destined for the external internet. I.e., hosts will act as if they are multi-homed.
3. A protocol is required for allocation and de-allocation of port numbers by the DNAT router.
4. The DNAT router should probably check the port number in the TCP/UDP header of each packet outbound to the external network to ensure that it has in fact been allocated to the local host that generated the packet.
5. The assumption that the source port set by the local host, when transmitting to a remote (external) system, will be used as the destination port in any reply packets sent by the remote system is obviously a requirement of this protocol. This requirement raises the general issue of "NAT-friendly" applications. That is, must remote applications be cognizant that they are communicating with a host behind a DNAT (or standard NAPT) router? For DNAT the answer is clearly "yes." However, the DNAT method of port number – IP address mapping, which requires no packet modification by the DNAT

router, enables a single, simple solution to make *all* applications DNAT-friendly. It must only be possible for the local host to specify the destination port number to be used by the remote application when sending packets back to the host on the stub network. This requirement allows a unique mapping for all packets inbound to the stub network via the DNAT router.

IPsec Across a DNAT Network

The two obstacles to implementing IPsec with NAT are easily overcome using DNAT instead, as will now be described.

The method for running IPsec across DNAT starts from the basic technique of DNAT, but replaces the port number mapping with SPI mapping. Specifically, the DNAT router allocates to each of the local hosts on its stub network blocks of non-overlapping ranges of SPI values. When a given host negotiates an incoming IPsec SA with a remote system on the global IP network, it ensures that the SPI assigned for this SA is selected from its allocated block of SPI values. Here, incoming SA is taken to mean an SA that terminates at the host for inbound packets (i.e., packets destined to the host). For outgoing SAs, the SPI is selected by the remote system, and is irrelevant to the DNAT router's actions. In the event of multiple levels of incoming SAs which terminate at the host, and which are associated with a single connection, only the uppermost level SA needs to use an SPI selected from the allocated block of SPI values. Recall that the SPI is stored in the IPsec protocol header of the associated IP packet. For the uppermost level of SA, the IPsec protocol header is always in the clear for any and all combinations of protocol and mode. Therefore, the DNAT router will always have read access to the SPI in the uppermost level SA associated with any incoming IP packet.

The simple "light-weight" tunnel described in the summary of DNAT above is still used for local routing between the DNAT router and local hosts involved in external communication. However, in the case of IPsec over DNAT, the DNAT router does not look at TCP or UDP port numbers (they may not even be readable in the case of IPsec with ESP). For outgoing packets, the DNAT router simply removes the light-weight tunnel header and forwards the remaining packet on its external network interface; the SPI in this outgoing packet is not recorded, and has no use in this scheme. For incoming packets, the DNAT router maintains a mapping between local IP address of all hosts and SPI value. When a packet arrives from the external, global IP network, the DNAT router examines the SPI in the packet's outermost IPsec protocol header. As noted, this header is always in the clear (the fact that there may be nested levels of IPsec headers below the outermost one is of no consequence to this mapping method). The SPI value in the IPsec header is then used to determine the local IP address of the local host; a light-weight tunnel header is constructed and prepended to the packet, then the packet is forwarded to the local host. The local host removes the tunnel header, and processes the packet as usual. The DNAT router never needs to modify any of the received packet's contents.

Even though TCP/UDP port numbers are not used in this method for address mapping by the DNAT router, each local host must still be allocated blocks of non-overlapping ranges of port numbers for source port assignment, as is done in DNAT without IPsec. This is because once a packet from a local host is received by a remote system (on the external network) and processed by IPsec, the original IP header, TCP/UDP header, and payload must conform to the standard protocols. In particular, it must not be possible for two or more local hosts on a DNAT LAN to be able to use the same source port number, since the IP packets that they construct all use the *same* source IP address. If use of the same source port number by multiple local hosts *were* allowed, port collision would result in the event that the multiple hosts attempted to communicate with the same application on the same remote server. That is, once IPsec input processing at the remote system is complete, there is no way to distinguish which arriving packets came from which DNAT local host, *except* for source port number. Thus, the protocol of using allocated blocks of port numbers for source port assignment by the local hosts on a DNAT LAN is maintained when IPsec is implemented on the LAN. However, source port is no longer used for the address mapping by the DNAT router.

The port numbers are also required for output IPsec processing at the remote system for packets destined to local hosts on the DNAT LAN. As briefly described above, parameters carried in each originating IP packet are used as selectors to determine what output IPsec processing should be applied on a packet-by-packet basis. If the remote system terminates IPsec

connections with more than one local host on a DNAT LAN, port number is, again, the only way to distinguish to which host each packet should be sent, since the source IP address for all such packets is the same (i.e., the IP address of the DNAT router's external interface). Note that in this case, destination port number is used, but its value is specified by the local host (see the comment on "NAT-friendly" applications in the summary of DNAT above). With proper filtering by selectors, including destination port number, each IP packet gets the correct output processing and IPsec protocol header, and is bound to the correct SPI. When these packets arrive at the DNAT router, the mapping of SPI to local host IP address proceeds as described above.

In short, forwarding/routing in the case of IPsec over DNAT looks just like DNAT, except that SPI is used instead of TCP or UDP port numbers. This technique eliminates the first problem of NAT with IPsec, since the DNAT router never modifies the contents of the packets it forwards, and requires read access only to the SPI in the outermost IPsec protocol header (which is always in the clear). Source ports are still used to disambiguate connections by the remote server.

The remaining problem is establishment of IPsec SAs. This problem is solved by: 1) modifying the name space for binding public keys to the identities of their owners; and 2) configuring the DNAT router to act as a Local Certificate Authority (LCA). In addition, it is assumed that the DNAT router is registered with a higher-level CA; i.e., it has holds a certificate in which its public key that is bound to its global IP address, and which is validated by the higher-level CA.

The modified name space must satisfy two requirements: 1) it must provide any local host on any stub network a globally-unique identifier that includes the external IP address of the DNAT router; and 2) for each local host, it must include the range of port numbers allocated to the local host by the DNAT router. The first requirement can be met by any of several alternative identifiers. For the purposes of illustration within this disclosure, we shall choose a name space consisting of the concatenation of the DNAT router's global IP address and the host's local IP address. This identifier is global by virtue of the uniqueness of the DNAT router's global IP address, and the uniqueness within the stub network of the host's local IP address. Any local host on the DNAT stub network that wishes to participate in IPsec begins by receiving its allocated block of port numbers and SPI values from the DNAT router. It must then provide a public/private key pair that it has generated, and have its public key certified by the LCA (configured on the DNAT router) in the same way that the DNAT router has its public key certified by a high-level CA. That is, the local host is issued a certificate by the LCA which contains a binding between its public key and the combination of its identifier and port number range, and which is validated by the LCA. Normal SA negotiation between a local host and remote computer on the global IP network proceeds as defined by the IPsec protocols.

The first requirement ensures the remote system that the local host has the right to use the global IP address of the DNAT router. The second requirement ensures the remote system that the local host has the right to use the range of port numbers; i.e., that the DNAT router in fact allocated those port numbers to the local host. It also provides a method for informing the remote system the range of port numbers that should be associated with the local host. As described above, port numbers are required to disambiguate IP packets from/to hosts on the DNAT LAN, for input/output IPsec processing.

Note that port number maintenance by the DNAT router includes both allocation and de-allocation, and can be fairly dynamic. Local hosts can request additional port numbers, and the DNAT router can render an allocated range invalid (de-allocation). If IPsec is implemented as well, additional certificates must be issued by the LCA for each allocation of port numbers to a local host. In addition, the DNAT router must maintain a list of all certificates issued to its local hosts, and ensure that the associated ports are never de-allocated as long as the certificates with bindings to these ports are still valid. Alternatively, if the DNAT router is allowed to de-allocate ports, it must revoke any certificates with bindings to these ports. Certificate revocation must include notification to all remote systems that have active SAs established with the local hosts whose certificates have been revoked. De-allocation and certificate revocation may be required, for example, when a local host has a system crash. Finally, in the event of a system crash on the DNAT router, the DNAT router, once it recovers, must be able to either recreate its certificate or gracefully revoke invalid certificates.

This method of authentication is not restricted to the form of the name space described above. For example, a combination of the DNAT router's global IP address and user email address (where the user is on the local host) would also work. The only constraint is that the DNAT router acting as an LCA must possess a valid certificate giving it the right to certify identifiers drawn from the chosen name space. It should also be noted that this method can be made work within the scheme disclosed by the authors in which Mobile IP is integrated with DNAT [9]. In this case, the mobile node's home agent is also its LCA. The name space identifier used for illustration in this disclosure is a concatenation of the home agent's global IP address and the mobile nodes local address on its home network. This information is available to the mobile node even while it is roaming (i.e., temporarily residing on a foreign network). Therefore no aspect of Mobile IP integrated with DNAT as described in the above referenced disclosure precludes the implementation of IPsec with DNAT as presented in this disclosure. That is, the method presented here also extends to IPsec within the context of Mobile IP, allowing a mobile node to maintain an IPsec-protected connection while it is roaming. The only requirement is that the mobile node's home network is managed as a DNAT stub network in which the mobile node resides as a local host when it is not roaming.

Finally, we note that the concept of using a modified name space to provide a unique identifier to a computer that lacks a globally unique IP address is not restricted to a design based upon the LCA. It also possible to define a global CA using a modified name space, and eliminate the need for the LCA. However, such a name space is insufficient for the DNAT environment, since it does not include port number, and hence the guarantee to the remote system that a local host has the right to use a specific port number. Aside from this shortcoming, the method of configuring the DNAT router to act as an LCA is proposed here in order to supply a completely described system for implementation. Also, since stub networks exist, and DNAT is a method for sharing global IP addresses within stub networks, the LCA approach described here provides an implementation path that would build upon an existing infrastructure, rather than requiring a new and undeveloped infrastructure. That is, if a DNAT system is built to solve the problem for which it was designed, IPsec could be made to work with it without requiring a new infrastructure to support a global CA with a modified name space.

References

[1] P. Srisuresh, and K. Egevang, "The IP Network Address Translator (NAT)," Internet Draft <draft-rfced-info-srisuresh-05.txt>, February 1998 (work in progress).
[2] M.S. Borella, D. Grabelsky, I. Sidhu, and B. Petry, "Distributed Network Address Translation," Internet Draft <draft-ietf-borella-dnat-aatn-00.txt>, April 1998 (work in progress)
[3] Stephen Kent, and Randall Atkinson, "Security Architecture for the Internet Protocol," Internet Draft <draft-ietf-ipsec-arch-sec-03.txt>, February 1998 (work in progress).
[4] Stephen Kent, and Randall Atkinson, "IP Encapsulating Security Payload (ESP)," Internet Draft <draft-ietf-ipsec-esp-v2-04.txt>, March 1998 (work in progress).
[5] Stephen Kent, and Randall Atkinson, "IP Authentication Header," Internet Draft <draft-ietf-ipsec-auth-header-05.txt>, March 1998 (work in progress).
[6] Douglas Maughan, Mark Schertler, Mark Schneider, and Jeff Turner, "Internet Security Association and Key Management Protocol (ISAKMP)," Internet Draft <draft-ietf-ipsed-isakmp-09.txt,.ps>, March 1998 (work in progress).
[7] H.K. Orman, "The OAKLEY Key Determination Protocol," Internet Draft <draft-ietf-oakley-02.txt>, July 1997.
[8] D. Harkins, and D. Carrel, "The Internet Key Exchange (IKE)," Internet Draft <draft-ietf-ipsec-isakmp-oakley-06.txt>, February 1998 (work in progress).
[9] M. Borella, D. Grabelsky, J. Mahler, and I. Sidhu, "Integrating Network Address Translation and Mobile IP," 3Com Invention Disclosure, 3Com docket number 1868.CS, submitted internally 4 June 1998.

This Page is Inserted by IFW Indexing and Scanning Operations and is not part of the Official Record

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- ☐ BLACK BORDERS
- ☐ IMAGE CUT OFF AT TOP, BOTTOM OR SIDES
- ☐ FADED TEXT OR DRAWING
- ☐ BLURRED OR ILLEGIBLE TEXT OR DRAWING
- ☐ SKEWED/SLANTED IMAGES
- ☐ COLOR OR BLACK AND WHITE PHOTOGRAPHS
- ☐ GRAY SCALE DOCUMENTS
- ☐ LINES OR MARKS ON ORIGINAL DOCUMENT
- ☑ REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY
- ☐ OTHER: _____

IMAGES ARE BEST AVAILABLE COPY.
As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.